

Snowflake Cortex Search

From Manual Search to Intelligent Retrieval

Your teams spend hours hunting through PDFs and semi-structured data, only to get incomplete answers. Generic LLMs don't understand your business. evolv helps you deploy Cortex Search to deliver fast, accurate results from your enterprise data – without the overhead of building a custom stack.

What's holding you back

- No attribution or explainability in results
- No custom embeddings or reranking
- Can't fine-tune on your glossary or data catalog
- Higher cost than standalone vector DBs
- Disparate ingestion pipelines and infrastructure
- “We don't trust the answers – it's just a black box”

What success looks like

- › Higher precision and recall in search results
- › Lower hallucinations and faster retrieval latency
- › Index freshness SLAs and scalable ingestion
- › Increased user adoption of search/chat workflows
- › Seamless integration with agents and analysts
- › Governance, permissions, and explainability built in

How evolv helps



evolv helps teams operationalize Cortex Search to **deliver enterprise-grade retrieval without the engineering overhead.**



We streamline ingestion, tune chunking, and **align search with your business glossary and governance needs.**



We don't just enable search – we make it **trustworthy, scalable, and ready for real decisions.**

Why evolv?

- **Search grounded in your business context** – not generic LLMs
- **Governance, attribution, and explainability from day one**
- **Optimized** ingestion, chunking, and index freshness
- **Deep Snowflake expertise** to reduce vendor sprawl
- **Focused on adoption, precision, and measurable ROI**